Naval Research Laboratory
Washington, DC 20375-5320

NRL/MR/5520--96-7842

STOW Traffic and Related Issues

GAM D. NGUYEN
STEPHEN G. BATSELL

Communications Systems Branch
Information Technology Division

April 16, 1996

DISCLAIMER NOTICE

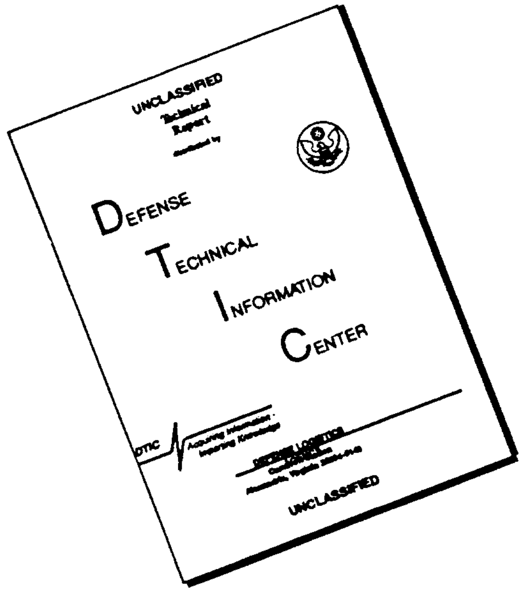

THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF PAGES WHICH DO NOT REPRODUCE LEGIBLY.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave Blank)*

2. REPORT DATE

April 16, 1996

3. REPORT TYPE AND DATES COVERED

Final

4. TITLE AND SUBTITLE

STOW Traffic and Related Issues

5. FUNDING NUMBERS

PE - 63226E
PR - B189/03 P5L10 2574
DPAM 5 0926
C - DABT63-92-C-0020

6. AUTHOR(S)

Gam D. Nguyen and Stephen G. Batsell

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER

NRL/MR/5520--96-7842

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Advanced Research Projects Agency (ARPA)
3701 North Fairfax Drive
Arlington, VA 22203-1714

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The Distribution Interactive Simulation (DIS) program has responded to the needs of having a common, standardized environment in which a large number of simulation nodes can interact in real time with each other. An early phase of the DIS program is the Synthetic Theater of War-Europe (STOW-E) demonstration, which was conducted in November 1994.

We analyze the STOW-E traffic using sample statistical techniques and find the following: At small time scales, the LAN traffic oscillates between burstiness and smoothness, whereas the WAN traffic exhibits smoothness. However, both the LAN and the WAN traffic exhibit high burstiness on larger time scales. Therefore, the traffic shows characteristics that would be produced by Poisson-type and self-similar processes. Additionally, in comparing the STOW-E traffic with an asymptotically self-similar process, we find that it has unusually high Hurst parameter (i.e., $H \sim 1$). It is well known that time-dependent processes can yield characteristics resembling correlated components or self-similarity. A simple mathematical model for the traffic is an uncorrelated process whose parameters vary according to another asymptotically self similar process of high burstiness. Since a combination of self-similar processes results in an asymptotically self-similar process, we conjecture that future STOW traffic should behave as asymptotically modulated self-similar processes.

14. SUBJECT TERMS

STOW	Long-range dependence	DIS
Self-similarity	Hurst parameter	
Short-range dependence	WSS	

15. NUMBER OF PAGES

23

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std 239-18
298-102

CONTENTS

STOW TRAFFIC AND RELATED ISSUES

1 Introduction

We analyze the unclassified traffic data of the Synthetic Theater of War-Europe (STOW-E)[1] distributed simulation demonstration, which is an early phase of the Distributed Interactive Simulation (DIS) program sponsored by Advanced Research Projects Agency (ARPA). In this paper, we also make a prediction about the possible traffic in future STOW networks. DIS forms a common, consistent, and distributed virtual environment. Different types of simulators can interact using unified standard DIS protocols, which can interoperate with future simulation systems. Current DIS applications focus on military combat simulations; however, commercial uses of DIS are expected in the future. The STOW-E exercise was conducted from 4 - 7 November 1994. It linked sixteen world-wide sites in the continental United States, Germany, and England, creating a virtual world in which real time systems could jointly operate in a variety of domains including land, sea, and air.

In 1983, Defense Advanced Research Projects Agency (DARPA) initiated the Distributed Simulation program with the Simulation Network (SIMNET) architecture. In this architecture, nodes received all information broadcasted by other nodes. However, the receiving nodes only processed relevant information. For a technical overview and history of SIMNET, see [Kana91].

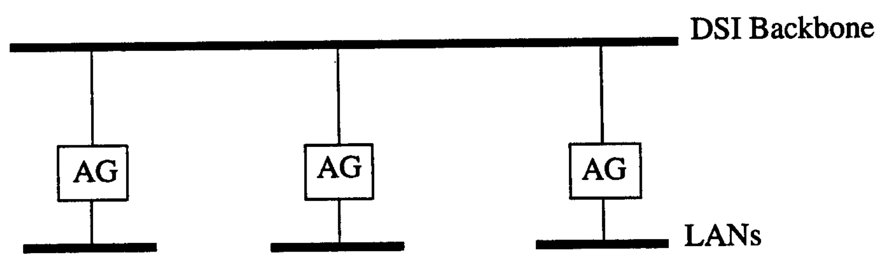

Fig. 1 STOW-E Network Configuration. Simulation LANs communicate with each other by way of the DSI WAN backbone and the AGs. The seven BRTs in the AGs reduce the unnecessary traffic on the backbone.

Figure 1 illustrates the conceptual configuration of the STOW-E network, which replaced the simple, yet inefficient, broadcast approach of SIMNET. The Defense Simulation Internet (DSI) WAN, composed of T1 links, functioned as the backbone for the STOW-E network. Simulation application hosts on ethernet LANs communicated to each other by way of the Application Gateways (AGs), which reduced the unnecessary traffic on the backbone. This was accomplished by using the following seven Bandwidth-Demand Reduction Technique (BRT) algorithms, which were housed in each AG: (1) *Protocol Data Unit (PDU) culling* to block the transmission of unnecessary PDUs over the WAN, (2) *grid filtering* to partition a terrain into regions for which updates are sent at different rates to accommodate different simulation fidelities, (3) *quiescent entity determination* to identify inactive entities and report them to all remote AGs, (4) *Protocol Independent*

1. Acronyms used in this paper are listed in Appendix B.

Manuscript approved March 20, 1996.

Compression Algorithm (*PICA*) to remove information redundancy by sending the bit-pattern differences between a reference and the actual bits, (5) *bundling* to combine PDUs into larger User Datagram Protocol/Internet Protocol (UDP/IP) packets to reduce packet rates, (6) *overload management* to spread packet transmissions out in time to reduce traffic burstiness, and (7) *LAN filter* to select and forward only relevant updates to each LAN. The BRT algorithms are explained in detail in [HCNF94]. As we will see in a later section, the seven BRT algorithms drastically shaped the traffic on the WAN. In STOW-E, the AGs bundled DIS PDUs into 1000-byte packets with a maximum time-to-fill interval of 0.04 second. The AGs examined and made complex decisions on whether and how to pass the LAN-generated-packets on to the DSI WAN. The only WAN traffic not traveling on any LAN was AG-to-AG control traffic.

The next step in the evolution of the STOW program is the STOW-97 system, in which High Performance AGs (HPAGs) will replace the STOW-E AGs. The HPAGs exploit new advances in networking technologies such as multicasting, multimedia services, and resource reservation. They decouple (interface) higher layer simulation applications from (and) the underlying network layer to facilitate independent developments among different system components. Besides HPAGs, STOW-97 will incorporate Agents that act as enablers for simulation applications. The STOW-97 architecture will implement legacy system translators to allow backward compatibility with older STOW systems such as SIMNET and STOW-E. A vision of the long-term STOW (i.e., the eventual Advanced Distributed Simulation or ADS) system architecture beyond the 1997 time frame is given in [CSTH95].

Recently, network researchers have shown increasing interests in self-similar processes. Simply speaking, a self-similar process exhibits burstiness over a wide range of time scales. In contrast, Poisson-type processes used in conventional traffic models become smooth after only one or two time scales. New studies report that many different types of traffic possess self-similarity, including ethernet [LTWW94] and Variable Bit Rate (VBR) video traffic [GaWi94]. Accurate characterizations of the traffic data are vital in any network design and management scheme. Therefore, one goal of this paper is to relate the STOW-E traffic to self-similar processes. As shown later, the traffic shows characteristics that resemble those of Poisson and self-similar processes.

In Section 2, we plot the overall series of the measured STOW-E traffic. To facilitate visual investigation of the traffic, different levels of data zooming are given in a variety of timing scales. Section 3 gives the mathematical framework of long-range dependent processes and of self-similarity. Section 4 applies several well-known graphical techniques such as pox and time-variance diagrams to determine quantitative characteristics of the traffic data. These plots suggest that the STOW-E traffic possesses both Short-Range Dependence (SRD) and Long-Range Dependence (LRD). The section also gives an overview of existing models for generating synthetic self-similar traces. We discuss implications from the analyzed STOW-E traffic on future STOW traffic in Sections 5. Section 6 concludes the paper.

2 STOW-E Traffic Data

In this paper, we analyze 4 of the 25 traces of the unclassified simulation traffic data from the Aviation Test Bed (AVTB) at Fort Rucker (Alabama, United States), which was one of two unencrypted sites of the STOW-E exercise. The other unencrypted site was the SIMNET simulator suite at Grafenwoehr, Germany. The DIS nodes at Fort Rucker operated at the allocated band-

width of 600 Kbits/s, and the SIMNET nodes at 1760 Kbits/s. In [NgBa96], we present a statistical compilation of all 25 STOW-E traffic traces, and find that the four traces presented here are representative.

Naval Command Control and Ocean Surveillance Center, Research, Development, Test and Evaluation Division (NRaD) designed and implemented the DIS Protocol Dlogger and the WANLogger for collecting the unclassified data used in our analysis. The Dlogger recorded all DIS traffic on a simulation LAN, whereas the WANLogger recorded the network traffic on the WAN side of the AG [NRaD95]. Each of the 25 data traces contains the number of traffic bytes[1] per 0.01 second. Data trace durations last from 4 minutes to more than 3 hours. In this paper, we discuss our analysis of the four traces shown in Table 1 for the following reasons. First, the first 2 traces contain the most numbers of data points for each WAN and LAN traffic category. Second, unlike 23 other traces, which contain LAN and WAN traffic at different times, the last two traces contain LAN and WAN traffic at the same time. Hence, one can consider the AG at Fort Rucker to be a system with LAN2 as input and WAN2 as output. The BRT algorithms, buffer sizes, software and hardware configurations in the AG determine the characteristics of the system.We devote all of the main sections in this paper to the analysis of the first two traces, and briefly discuss the results for the last two traces in Appendix A. In Table 1, the column Null-Interval Count records the number of 0.01-second intervals that have no traffic (i.e., zero-byte values). The Sparseness column gives the ratio of null intervals to sample sizes; i.e., Sparseness = Null-Interval Count/Sample Size.

TABLE 1. Sample Statistics for STOW-E Traffic

Traces	Time Duration	Sample Size	Null-Interval Count	Sparseness	Sample Mean	Standard Deviation
WAN1	3:42:14	1333416	1221432	.92	81	275
LAN1	2:04:51	749139	628127	.84	247	903
WAN2	0:30:00	180086	162207	.90	15	47
LAN2	0:30:00	180038	172118	.96	15	83

As the first step in the analysis, we plot the entire time series for trace LAN1 and trace WAN1, in the finest time units of 0.01 second, in Figs. 2(a) and 3(a), respectively. Notice a marked contrast in the traffic profiles in Figs. 2(a) and 3(a): the WAN traffic is much smoother than the LAN traffic as a result of the BRT algorithms in the AG at Fort Rucker.

To further investigate the burstiness of the traffic, following [LTWW94], we plot the data on several different larger time scales ranging from 0.1 to 10 seconds in Figs. 2 and 3. In Fig. 3, the WAN traffic looks smooth on the first time scale (0.01 second) and then becomes bursty on larger scales (0.1 - 10 seconds). In contrast, the LAN traffic in Fig. 2 exhibits burstiness on all four time scales. We will quantify the level of the traffic burstiness in Section 4. The above observations are consistent with the fact that the BRTs are effective, at least on the first time scale in smoothing out

1. The DSI WAN and the Ethernet LANs use broadcasting network protocols; therefore, the data points represent both the transmitted and received traffic.

the traffic before it is injected into the WAN backbone. Using the Bundling algorithm, each AG collects LAN PDUs and bundles them into larger packets for transmission over the WAN. A larger packet accumulates the PDUs based on a timer and a buffer limit. Then the bundling packet is transmitted to the WAN. Additionally, by setting up a maximum upper limit for the WAN-bound packet rate, the Load Leveling algorithm smooths out short-term traffic spikes.

To study the STOW-E traffic in detail, we further plot both the LAN and the WAN traces for disjoint time windows in Figs. 4 - 9. Each of the plots is composed of 1000 samples, and the sample size varies from 0.01 second to 1 second, as shown in the figures. Notice that the plots of scales 0.01 and 0.1 second in Figs. 4 and 5 show that the LAN traffic behaves like a periodic train of clusters. Each cluster has its own probability distribution. The start of each cluster is associated with data initiated from user simulation software, whereas pulses within each cluster correspond to data generated by network hardware or software. This observation agrees with the train model introduced in [JaRo86].

3 Long-Range Dependence Processes and Self-Similarity

In this section, we review the mathematical framework of long-range dependence and of self-similarity. Let X_n be a discrete-time process representing the number of traffic bytes collected at time n (i.e., the number of bytes collected in the nth 0.01-second time unit). We also assume that X_n is wide-sense stationary (WSS). Let $X_k^{(m)}$ be a WSS process formed by averaging the process X_n in nonoverlapping blocks of m, i.e.,

$$X_k^{(m)} = (X_{km-m+1} + \dots + X_{km})/m. \tag{1}$$

The variance of $X_k^{(m)}$ satisfies the Yule equation:

$$\operatorname{var} X_k^{(m)} = (\operatorname{var} X)/m + (2/m^2)\sum_{s=1}^{m-1}\sum_{h=1}^{s} r(h), \tag{2}$$

where $r(n) = E\{(X_{n+1}-m)(X_1-m)\}$ is the autocorrelation of the underlying process. There are a number of interesting consequences of (2). If $r(i) = 0$ (e.g., X_n is uncorrelated), $\operatorname{var} X^{(m)} = (\operatorname{var} X)/m$; therefore, the variance of the aggregated process decays hyperbolically with the aggregated size m. More generally, if

$$\sum_{h=1}^{\infty} r(h) < \infty, \tag{3}$$

the variance of the aggregated process decays hyperbolically with m. Therefore, a WSS process whose autocorrelation satisfies (3) is defined to be short-range dependent (SRD) [Cox84].

In contrast to the SRD process whose aggregated process must have hyperbolic-decaying variance, a process is long-range dependent (LRD) when $\sum_{h=1}^{\infty} r(h) = \infty$. In particular, a LRD process with autocorrelation function of the form $r(i) \propto i^{-\beta}$, $0 < \beta < 1$ is called self-similar (see [Cox84] and [LTWW94]), where $\propto$ denotes proportionality. From this simple definition, it can be shown that a WSS process $\{X_n\}$ is self-similar if its averages defined in (1) satisfy $\operatorname{var} X^{(m)} \propto m^{-\beta}$, for all m. Power spectral density is the Fourier transform of the autocorrelation and, for a self-similar process with parameter β, it has the form [Cox84]

$$g(f) = \frac{1}{2\pi}\sum_{k=-\infty}^{\infty} r(k)\, e^{-j2\pi fk} \propto f^{-(1-\beta)} \text{ as } f \to 0. \tag{4}$$

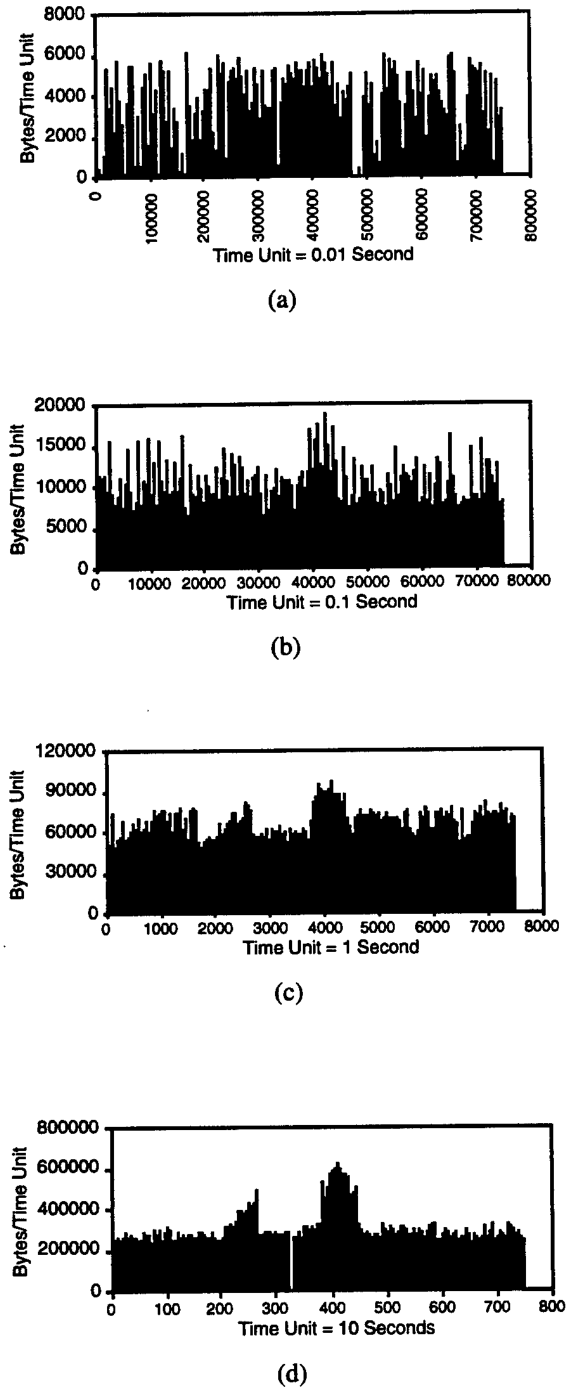

Fig. 2 STOW-E Traffic Trace LAN1 at 4 Time Scales: 0.01 - 10 Seconds. The traffic shows burstiness in all four time scales.

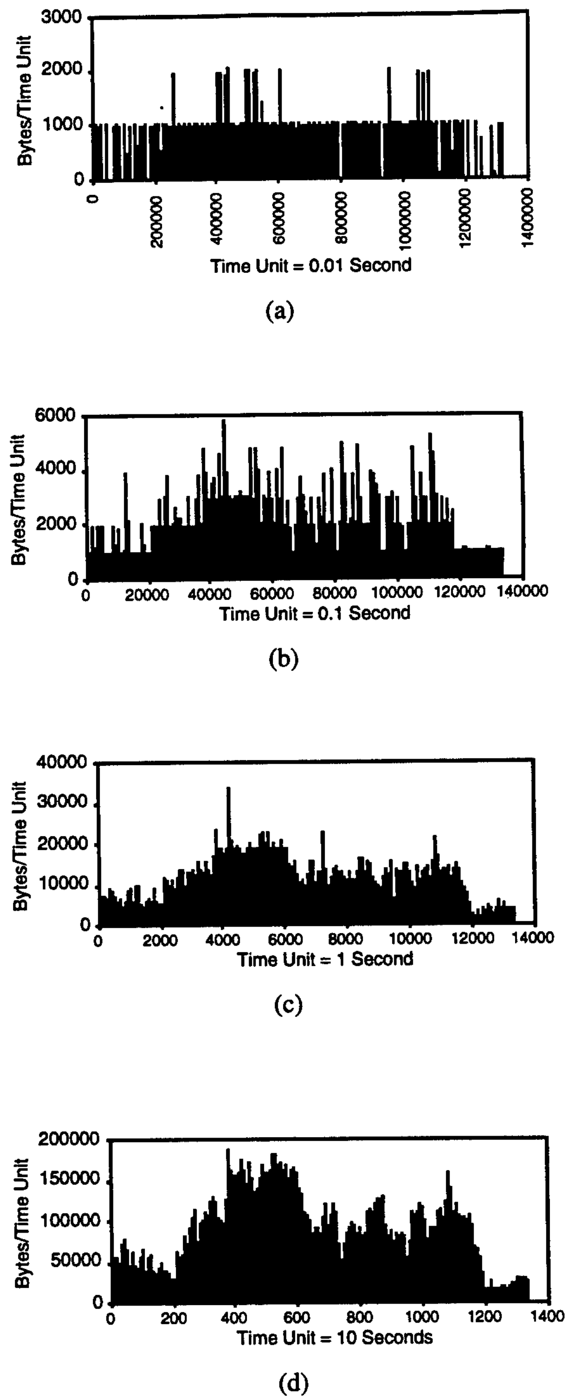

Fig. 3 STOW-E Traffic Trace WAN1 at 4 Different Scales: 0.01 - 10 Seconds. The BRT algorithms housed in the AG smoothen the WAN traffic in the 0.01-second time scale.

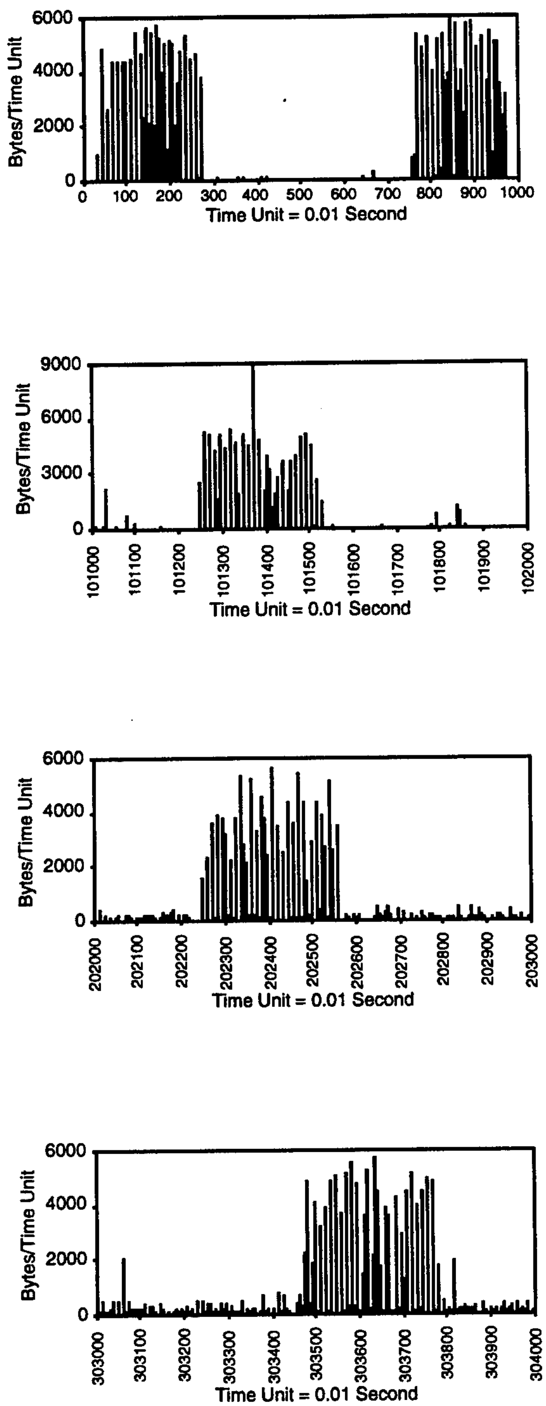

Fig. 4 Traffic Trace LAN1 in Four Non-Overlapping Time Windows. Each window contains 1000 0.01-second samples. The traffic forms clusters of approximately equal magnitudes (~5000 bytes).

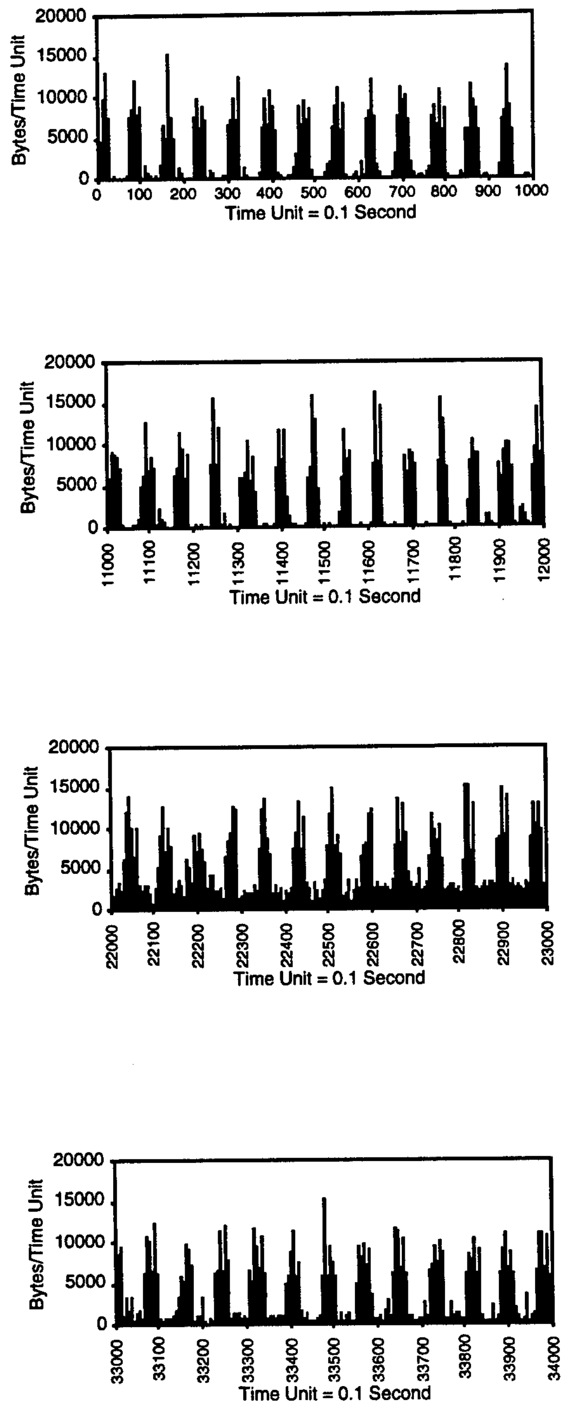

Fig. 5 Traffic Trace LAN1 in Four Non-Overlapping Time Windows. Each window contains 1000 0.1-second samples. The traffic forms clusters of approximately equal magnitudes (~10000 bytes).

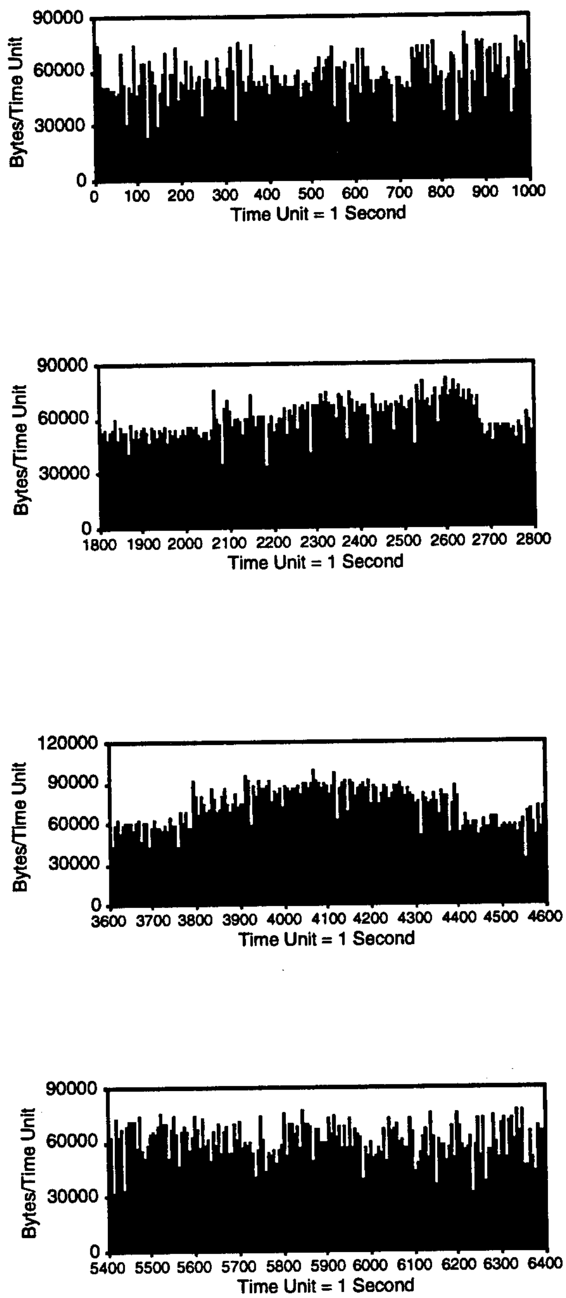

Fig. 6 Traffic Trace LAN1 in Four Non-Overlapping Time Windows. Each window contains 1000 1-second samples. The traffic remains bursty in all four windows.

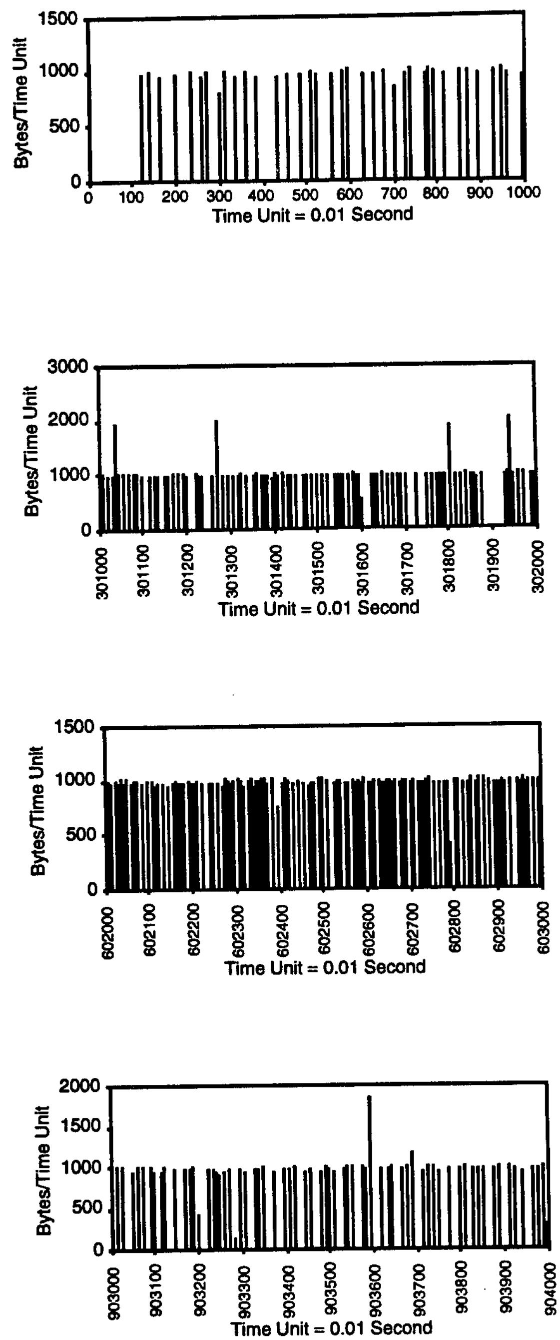

Fig. 7 Traffic Trace WAN1 in Four Non-Overlapping Time Windows. Each window contains 1000 0.01-second samples. The WAN traffic is more uniform than the LAN traffic as a result of the seven BRT algorithms housed in the AG (see Fig. 4). WAN traffic-load level is about 1000 bytes per 0.01 second.

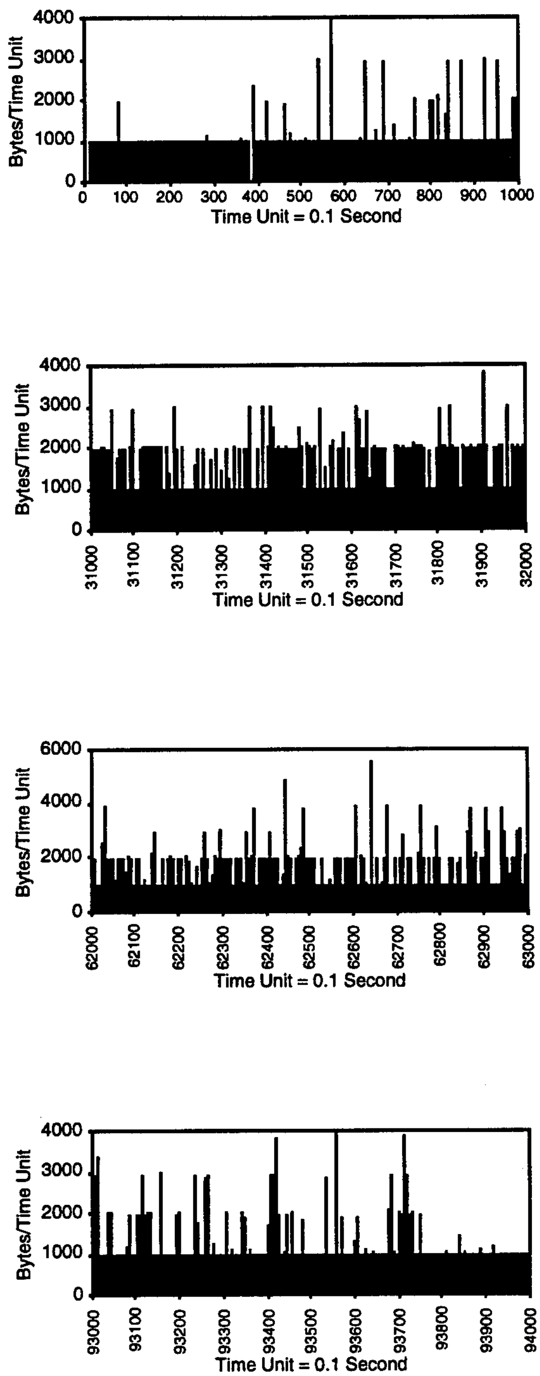

Fig. 8 Traffic Trace WAN1 in Four Non-Overlapping Time Windows. Each window contains 1000 0.1-second samples. Traffic-load level is no longer uniform, which suggests either asymptotically LRD or time-dependence of the WAN traffic.

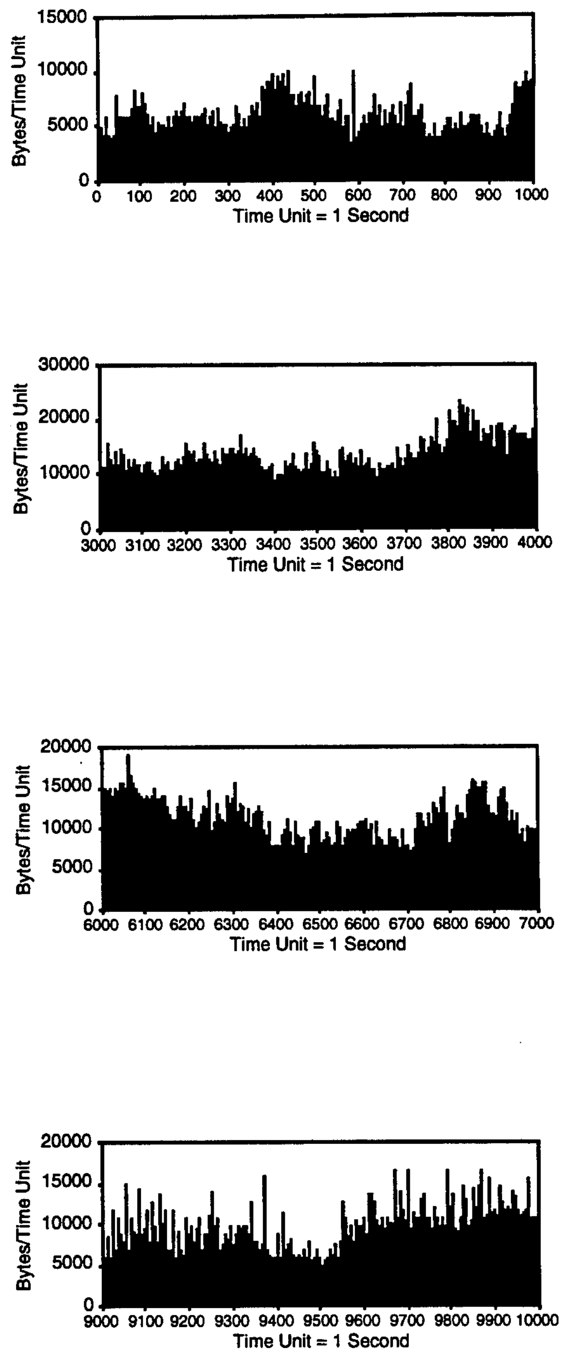

Fig. 9 Traffic Trace WAN1 in Four Non-Overlapping Time Windows. Each window contains 1000 1-second samples. The traffic remains bursty in all four windows.

A WSS process can also be classified as LRD or SRD by examining its Rescaled Adjusted Range or the R/S statistic. Let $\bar{X}$, $S^2(n)$ be the sample mean and variance of the observation WSS process X_n, respectively, and let $W_k = \sum_{i=1}^{k} X_i - k\bar{X}$, $k = 1, 2, \ldots, n$. The rescaled adjusted range is defined by

$$R(n)/S(n) = (\max(0, W_1, W_2, \ldots, W_n) - \min(0, W_1, W_2, \ldots, W_n))/S(n). \qquad (5)$$

Then in many cases of significant theoretical and empirical interests, $E[R(n)/S(n)] \propto n^H$, as $n \to \infty$, where $H = 1 - \beta/2$ is the Hurst parameter. H is a measure for burstiness; i.e., the higher H the burstier the process [LTWW94]. The underlying WSS process is of SRD when $H = 0.5$ and of LRD otherwise.

In [HDLK95], Huang et al. show that the steady-state behavior of a process formed by multiplexing two heterogenous self-similar sources will be that of the burstier one. Therefore, self-similarity is both persistent and dominating in the long run.

4 Graphical Sample Statistics for STOW-E Traffic

Recall from Section 3 that an asymptotic self-similar process $\{X_n\}$ of Hurst parameter $H = 1 - \beta/2$, $0 < \beta < 1$, has the following properties [Cox84]:

(4.0) $\{X_n\}$ is WSS;

(4.1) The autocorrelation $r(i) \propto i^{-\beta}$ as $i \to \infty$;

(4.2) The power spectral density $g(f) \propto f^{-(1-\beta)}$ as $f \to 0$;

(4.3) The variance of the aggregated process var $X^{(m)} \propto m^{-\beta}$ as $m \to \infty$;

(4.4) The rescaled adjusted range $E[R(n)/S(n)] \propto n^H$ as $n \to \infty$.

In this section, we use available graphical techniques to test the self-similar characteristics of the STOW-E traffic [LTWW94]. Notice that all the moments in (4.1) — (4.4) are computed by using sample estimates; therefore, inferences for LRD or self-similarity only refer to features produced by the use of sample statistics. In other words, in this analysis, LRD or self-similarity should be interpreted as *sample* LRD or *sample* self-similarity. Then the unknown parameters H and β can be estimated by plotting (4.1) — (4.4) in log scales[1]. Figure 10 depicts a *R/S* plot (also called *pox* plot), a variance-time plot, a power spectral density plot (using periodogram estimate), and an autocorrelation plot for trace LAN1, respectively. Figure 11 gives the corresponding plots for trace WAN1.

In each pox plot, the straight line of slope 0.5 corresponds to an uncorrelated (i.e., non-bursty) process; the line of slope 1 corresponds to a process with a very high degree of burstiness. In each variance-time plot, the line of slope -1 or 0 corresponds to an uncorrelated process or highly bursty process, respectively. Notice that each pox plot in fact conforms with a multi-valued function of the lag variable d [LTWW94][2]; therefore, it is useful only if it exhibits a consistent degree of the sought characteristics. In this regard, if the self-similarity manifestation of the data is not

1. From (4.4), $\log E[(R(n))/S(n)] \sim H \log n$; therefore, H is the slope of the *R/S* plot. Similarly, β is the slope of the variance-time plot.

2. In computing the *R/S* statistics, the n samples are segmented into windows of equal size d. Therefore, more *R/S* statistics are produced for smaller values of d.

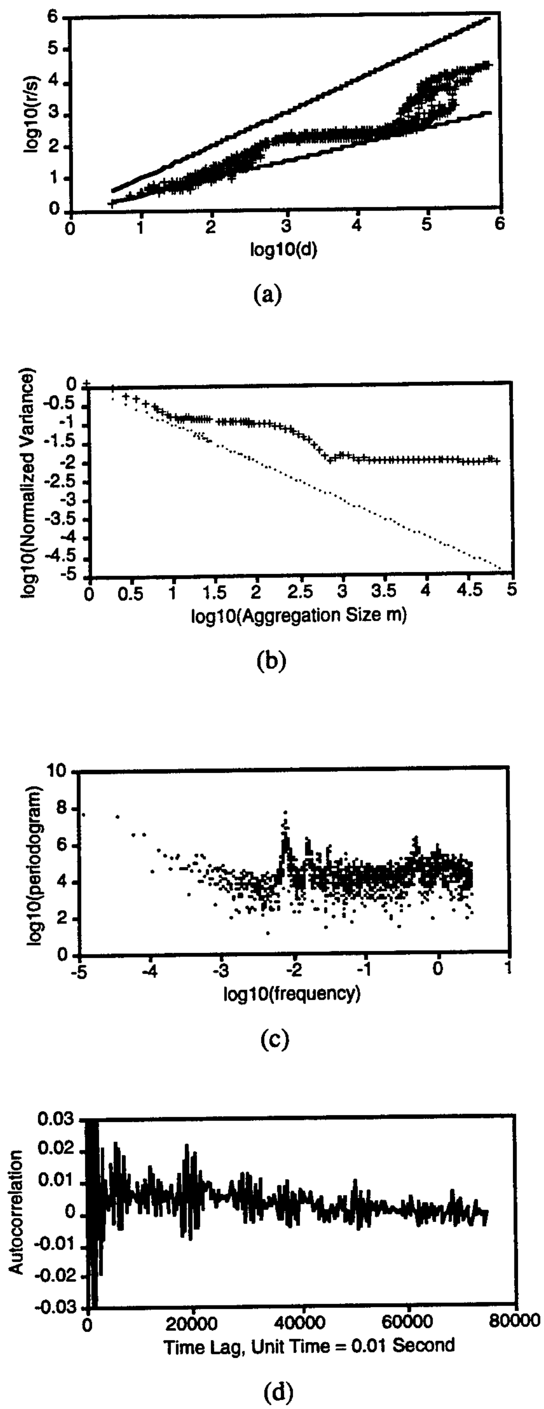

Fig. 10 Graphical Statistics for Trace LAN1. (a) *R/S* Plot: *d* denotes lag values or window sizes, which are different from the aggregation sizes *m*. The reference straight line of slope 1 (0.5 resp.) corresponds to H=1 (0.5 resp.). (b) Variance-Time Plot: The reference straight line of slope -1 corresponds to H=0.5. (c) Periodogram Plot. (d) Autocorrelation Plot.

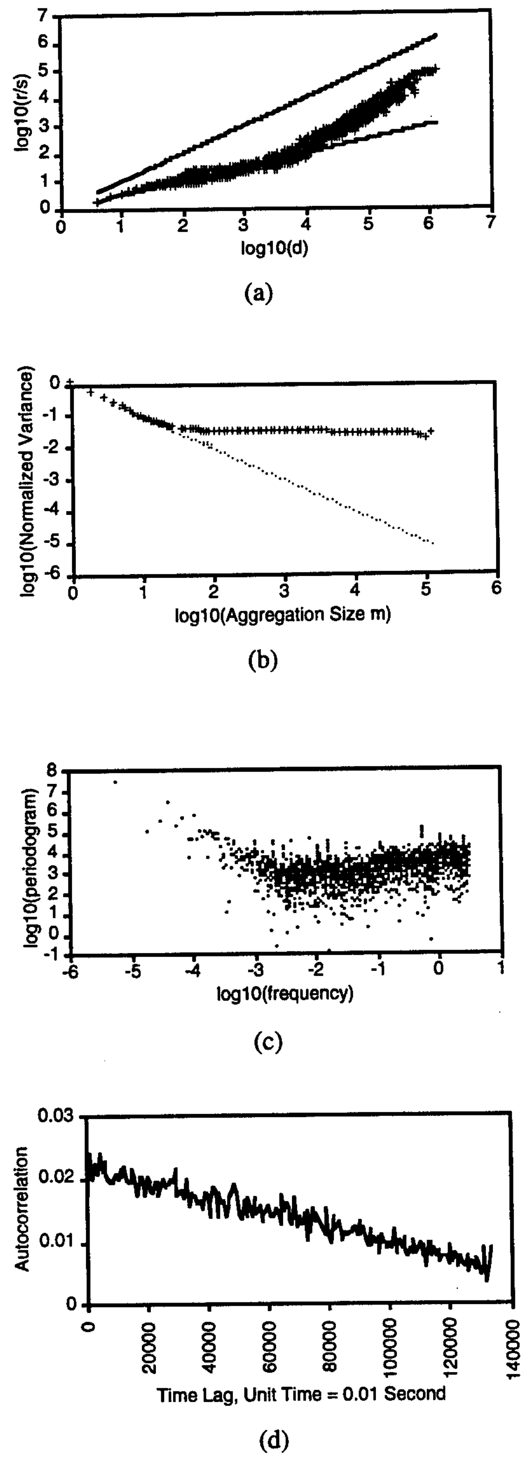

Fig. 11 Graphical Statistics for Trace WAN1. (a) *R/S* Plot: *d* denotes lag values or window sizes. The reference straight line of slope 1 (0.5 resp.) corresponds to H=1 (0.5 resp.). (b) Variance-Time Plot: The reference straight line of slope -1 corresponds to H=0.5. (c) Periodogram Plot. (d) Autocorrelation Plot.

evident from the pox plots alone, other techniques such as variance-time plots must be used instead. The WAN traffic, as a function of the aggregated size m, starts out as a non-bursty process as manifested in both the pox and time-variance plots in Figs. 11(a) and 11(b), then becomes sharply more bursty as m increases. Therefore, it has a flavor of an asymptotically self-similar process. From the pox plot in Fig. 10(a) the LAN traffic, as a function of the aggregated size m, starts out non-bursty, becomes more bursty with increasing m, then turns to a less bursty state, and turns to another bursty state when m grows further. It is unclear from the pox plot alone if the LAN process goes to a less bursty state or continues its trend when m continues to increase still further. However, by inspecting the plots over different scales in Fig. 2 as well as the variance-time plot for the LAN trace in Fig. 10(b), it still exhibits high degree of burstiness at larger scales; thus, the LAN traffic is comparable to an asymptotically self-similar process. However, it is important to keep in mind the fact that both the LAN and the WAN traffic also possesses uncorrelated (i.e., Poisson-type) components.

The asymptotic slope of the WAN pox plot (Fig. 11(a)) is unusually high (i.e., more or less close to unity), whereas that of the LAN traffic is inconclusive. On the other hand, the variance-time plots for both the LAN and WAN trace possess an asymptotic slope of about zero, which corresponds to a unity Hurst parameter. Examining both the periodogram and autocorrelation plots for the traffic traces also suggests that the Hurst parameter is high, see (4.1) and (4.2). Because higher H means higher burstiness, the STOW-E traffic exhibits an unusually high degree of burstiness in large time scales. Recall that self-similarity is persistent and that aggregation of self-similar processes will be governed by the process with the highest degree of self-similarity. Therefore, when other types of traffic and STOW-E traffic are integrated into a heterogenous environment, the resulting traffic should have characteristics of the asymptotically bursty STOW-E traffic.

The Hurst parameter for the STOW-E traffic departs from 0.5 in large time scales because of at least two possible reasons. First, the traffic is asymptotically self-similar if it fulfills WSS[1]. The other reason is the possibility that the STOW-E traffic has parameters that vary with time. Inspection of Figs. 2 - 8 suggests that this is the case. Recall that the traffic also displays uncorrelated characteristics in small time scales. Therefore, for all time scales, we model the STOW-E traffic as *an uncorrelated process over non-overlapping time intervals with an approximately constant parameter (e.g., traffic load) in each interval; however, the parameter will vary from one time interval to the next. Its numerical value will determine the traffic load in each time interval* (cf. [SlLe95]).

Plots of (empirical) complementary cumulative distribution functions for the WAN and LAN traffic in log scales are shown in Figs. 12 and 13, respectively. Recall that the complementary cumulative distribution function is one minus the cumulative distribution function. Notice that there are deviations from the abscissas (i.e., down jumps) at the beginnings for both plots because STOW-E messages had a minimum length (~144 bytes for LAN and ~80 bytes for WAN) and the traffic traces contain many null values (see Table 1). The plots do not fit well to any well-known distributions such as Normal, Gamma, Lognormal, or Pareto. However, the tails of the plots behave somewhat as straight lines, hence they match the heavy-tails property of Pareto distributions (cf. [GaWi94]).

1. By definition, self-similarity implies WSS. Section 5 discusses WSS and related issues.

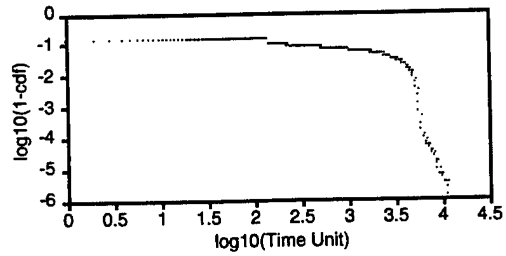

Fig. 12 Complementary Cumulative Distribution of Trace WAN1. The right tail varies somewhat as a straight line, which agrees with Pareto characteristics.

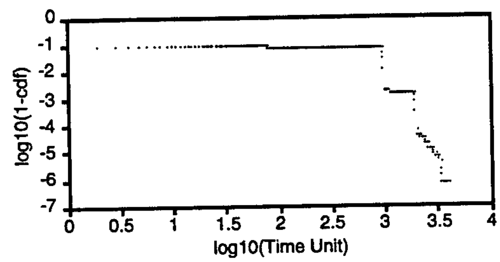

Fig. 13 Complementary Cumulative Distribution of Trace LAN1.The right tail varies somewhat as a straight line, which agree with Pareto characteristics.

Several methods exist for generating synthetic traces of exactly or pseudo self-similar traffic [Paxs95]. From the pox and variance plots (Figs. 10 and 11), the STOW-E traffic behaves like a process modulated by an asymptotic self-similar process. The STOW-E traffic for both the WAN and LAN possesses a certain degree of short-term effects, which cause a deviation from exact self-similarity. A simple model for an asymptotically self-similar process is a linear combination of exactly or asymptotically self-similar processes. The resulting process, as a function of the aggregated size *m*, will start out as a transient process then will become self-similar as *m* grows larger.

A simple way to generate an asymptotically self-similar process, at least in principle, uses the $M/G/\infty$ queue model with Pareto service times [Cox84]. Specifically, let the queue input be a Poisson process with an arrival rate ρ, and the service times be Pareto random variables with the identically independent distribution $P\{\text{Service Time} \le x\} = 1-(a/x)^{\beta}, a, \beta \ge 0, x \ge a$.Then the queue output process obeys asymptotically self-similarity with autocorrelation function

$$r(k) = [\rho a^{\beta}/(\beta-1)]\,k^{(1-\beta)}. \tag{6}$$

Furthermore, the self-similar output process has Poisson marginal distribution with mean $\rho\beta a/(\beta-1)$.

Another method of producing an asymptotically self-similar process is Hosking Fractional Differencing. Formally, let

$$\nabla^{d} X_n = A_n, \tag{7}$$

where A_n are independent identically distributed random variables, $d = (1-\beta)/2 = H - 1/2$ and

$$\nabla^d = (1-B)^d = \sum_{k=0}^{\infty} \binom{d}{k} (-B)^k, \tag{8}$$

with B denoting the backward shift operator defined by $BX_k = X_{k-1}$, and all factorials are expressed in terms of Gamma functions. The autocorrelation function of X_n satisfies

$$r(k) = \frac{d(1+d)\ldots(k-1+d)}{(1-d)(2-d)\ldots(k-d)}, \tag{9}$$

$$\text{as } k \to \infty, \quad r(k) \sim \frac{(-d)!k^{2d-1}}{(d-1)!}. \tag{10}$$

An algorithm, with A_0 being a Gaussian random variable of mean 0 and variance v_0, for generating the Hosking Fractional Differencing process is also available [GaWi94].

5 Implications

The BRT algorithms in the AG make the WAN traffic more uniform on smaller scales (less than 0.1 second). However, on greater time scales, the WAN traffic remains as bursty as the preprocessed LAN traffic. This shows that the BRT algorithms housed in each AG are effective in reducing the higher frequency components of the traffic; nevertheless, they have not reduced the low frequency components of the traffic. This observation agrees with the currently prevailing view that self-similar-type traffic poses a challenge for network congestion management. In particular, for networks of self-similar traffic, the overall packet-loss rate decreases very slowly with increasing buffer sizes, which in turn increase packet delays.

The impacts of self-similarity or long-range dependence on present and future networks deserve more study. In [PaFl94], Paxson and Floyd observe that Poisson models, or models that do not accurately reflect the long-range dependence, underestimate performance measures such as average packet delays or maximum buffer sizes. If the traffic does obey the Poisson models, network congestion that causes packet drops can be easily handled by linearly increasing the buffer sizes. In contrast, long-range dependent traffic can cause very long periods of congestions that can not be relieved by the methods used for the Poisson models. Furthermore, application hosts producing self-similar traffic, if granted higher priority, will starve the lower priority hosts for long periods of time. Self-similar traffic can also have a deleterious impact on the admission control polices that are based on recent traffic levels. Such policies, even though effective for certain types of traffic, do not work well for highly bursty traffic, which can produce a long peak period right after the measured low load period upon which the policy is based.

Having mentioned the implications of truly LRD or self-similarity on network management, one still has to face with the possibility of non-WSS traffic. By definition, LRD or self-similar processes must be WSS. Furthermore, generally, sample analysis used in this paper is accurate only for WSS series. Are the empirical STOW-E traces, as well as traces collected by other researchers, reasonably approximated by WSS processes? Mathematically, WSS is a simple and a precise concept. A process $\{X_n\}$ is WSS if it satisfies the following three conditions:

(5.1) $E(X_n^2) < \infty$ for all n,

(5.2) $E(X_n) = m$ for all n, and

(5.3) $E\{(X_r - m)(X_s - m)\} = E\{(X_{r+t} - m)(X_{s+t} - m)\}$ for all r, s, and t.

Suppose that one has collected N empirical data points $X_1, X_2, \ldots, X_N$, where N is a "large" number, say $N = 10^5$. How can one test to see if the data is WSS? At first sight, the problem seems straightforward. Define $S_n = (X_1 + X_2 + \ldots + X_n)/n$ as the sample mean, and test to see if S_n approaches to a constant when n increases. Suppose the answer is yes; i.e., $S_n \to \mu$ for some finite number μ. Can one conclude that $EX_n \approx \mu$ as required by the condition (5.2)? A simple counter-example is the process $X_n = 1/n$. Clearly, the process is not WSS, even though its sample mean converges to zero; i.e., $\mu = 0$. Testing condition (5.3) is even more difficult if not impossible. Thus, validation of WSS for an empirical data set needs further investigations.

Therefore, in general, we have to assume that the data represent a non-WSS process for which sample analysis is no longer accurate. Many researchers suggest decomposing a non-WSS process X_n into two components

$$X_n = t_n + Y_n, \tag{11}$$

where t_n is a deterministic function (i.e., the trend component) that only depends on the time n and Y_n a WSS process. However, we have to identify the trend component in addition to testing to see if the second component is WSS. The problematic nature of WSS assumptions is discussed in detail in [Klem74] and [DLOR95].

As previously mentioned, many papers report evidence of self-similarity in networking traffic. Queueing analyses indicate that self-similar traffic will cause severe networking problems such as long packet delay and high packet loss. Have such problems occurred frequently in current networks? We underestimate the burstiness of the traffic by using the simple Poisson models. Do we overestimate the traffic burstiness by the self-similar models? Is the cost of underestimating far exceeding that of overestimating? What is the cost of assuming WSS (a necessary condition for self-similarity) when in fact the data do not represent a WSS process? As with other researchers, in this analysis, we consider the STOW-E empirical data as if they obey WSS so that we can use available sample statistical techniques in our studies. The STOW-E traffic, as seen from time-variance and pox plots, exhibits interleaving of uncorrelated and strongly correlated components. Should we pay attention to intermittent short term effects or only to long term effects? It is well known that time-dependent processes (i.e., non-WSS processes) can produce characteristics resembling correlated components or self-similarity [Klem74].

Other issues when analyzing the empirical data are the sizes of the traces and the rate of convergence of the results. Should we consider a long trace as a whole with many time varying parameters or should we segment the trace into shorter traces and study each one independently? Sample analysis shows that the STOW-E traffic is asymptotically bursty. What is the difference, concerning network management, between two processes which both have the same degree of burstiness eventually, but reaching the same level of burstiness at different time scales (i.e., with different rates). For example, one process has Hurst parameter 0.9 after scaling 1 second or more; and the other process has the same Hurst parameter only after scaling 10 seconds or more. Should we remove traffic irregularities such as system errors and malfunctions from our data before we actually do the analysis? Such irregularities always exist in real networks and network planners must take them into consideration when they design the network protocols.

6 Conclusions

The STOW-E traffic shows characteristics that would be produced by Poisson-type and self-similar processes. Since time-dependent processes can yield characteristics resembling strongly correlated components or self-similarity, the traffic can be modeled as an uncorrelated process modulated by another asymptotically self-similar process. Furthermore, the traffic exhibits a very high degree of burstiness in larger time scales. Popular traffic models include self-similar (e.g., ethernet and VBR video) and Poisson processes (e.g., user-initiated TCP session arrivals [PaFl94]). With the advent of high speed networking technologies, which in turn stimulate newer applications, more exotic types of traffic are possible in the future. However, traffic of sample self-similar types will play a major role for many years to come. Aggregation of self-similar processes, or aggregation of self-similar processes and Poisson-type processes, will produce a self-similar process with the degree of self-similarity determined by the component process of the highest self-similarity. Future STOW networks such as STOW-97 and the eventual ADS will support traffic of STOW-E type in addition to others; therefore, the resulting integrated traffic will be at least as bursty as STOW-E-like traffic.

Appendix A

This appendix briefly discusses the last two data traces in Table 1. The traces, which were measured at the same time (a 30-minute duration), contain the STOW-E traffic on each side of the AG at Fort Rucker; i.e., trace WAN2 represents the WAN traffic, whereas trace LAN2 represents the LAN traffic. We plot the overall traffic profiles for the WAN and the LAN trace in Figs. 14 and 15, respectively, at 3 different time scales since both traces only have about 180000 data points (i.e., relatively short traces). The LAN traffic exhibits high burstiness in all three time scales; whereas the WAN trace in the finest 0.01-second time scale looks smooth as a result of the BRT algorithms housed in the AG. However, high burstiness begins to appear when the aggregation size increases. This observation agrees with that for traces WAN1 and LAN1; i.e., the BRT techniques are only effective for short periods (i.e., for time scales of 0.1 second or less) in smoothing out the LAN traffic before it goes into the WAN. The WAN traffic remains as bursty as the LAN traffic in higher time scales.

The window plots in 0.01-second and 0.1-second units for trace LAN2 appear in Figs. 16 and 17, respectively. The traffic behaves like a train of clusters whose denseness varies according to the traffic-load level and time. Fig. 18 depicts the traffic windows for trace WAN2; the windows show similar characteristics to those of trace WAN1 in Fig. 7. Figures 19 and 20 give sample statistical quantifications for the two traces. Generally, we can draw the same conclusion as we did in Section 4 for traces WAN1 and LAN1; i.e., the LAN traffic possesses interleaving of smoothness and burstiness in small time scales and becomes more bursty in larger scales, whereas the WAN traffic shows a more consistent degree of smoothness in small time scales and high degree of burstiness in larger time scales. Therefore, a simple mathematical model for the traffic is an uncorrelated process modulated by another asymptotically self-similar process. The last two figures (Figs. 21 and 22) depict the empirical complementary cumulative distribution functions for the traces. As for traces WAN1 and LAN1, the distribution functions for traces WAN2 and LAN2 resemble somewhat as straight lines at the far right ends, consistent with the heavy-tail characteristics of Pareto distributions.

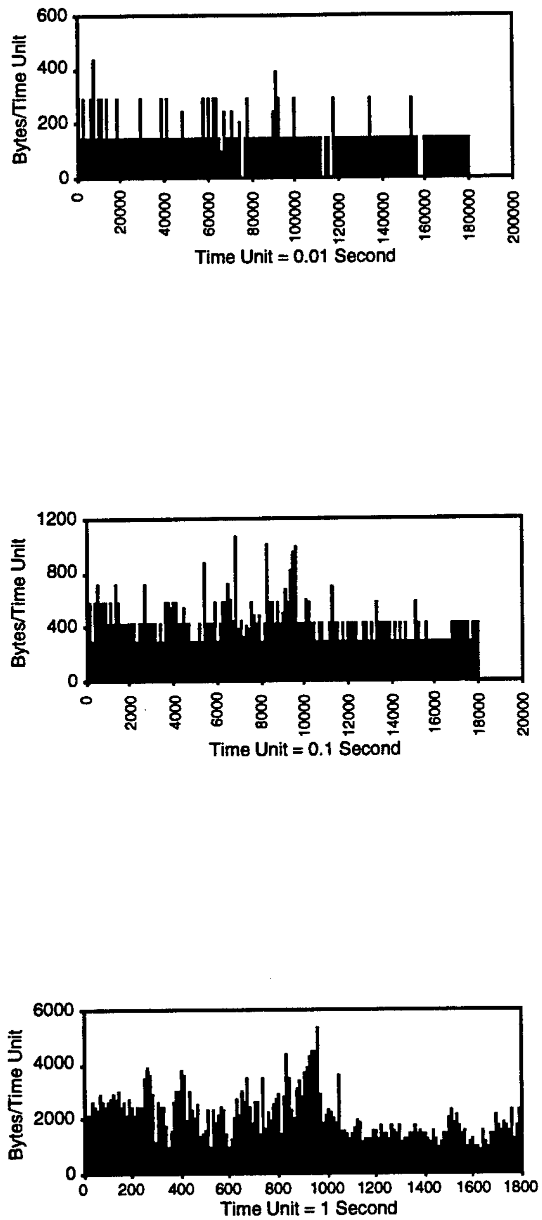

Fig. 14 Trace WAN2 at 3 Different Time Scales: 0.01 - 1 Second. The BRT algorithms housed in the AG smoothen the WAN trace in the 0.01-second time scale.

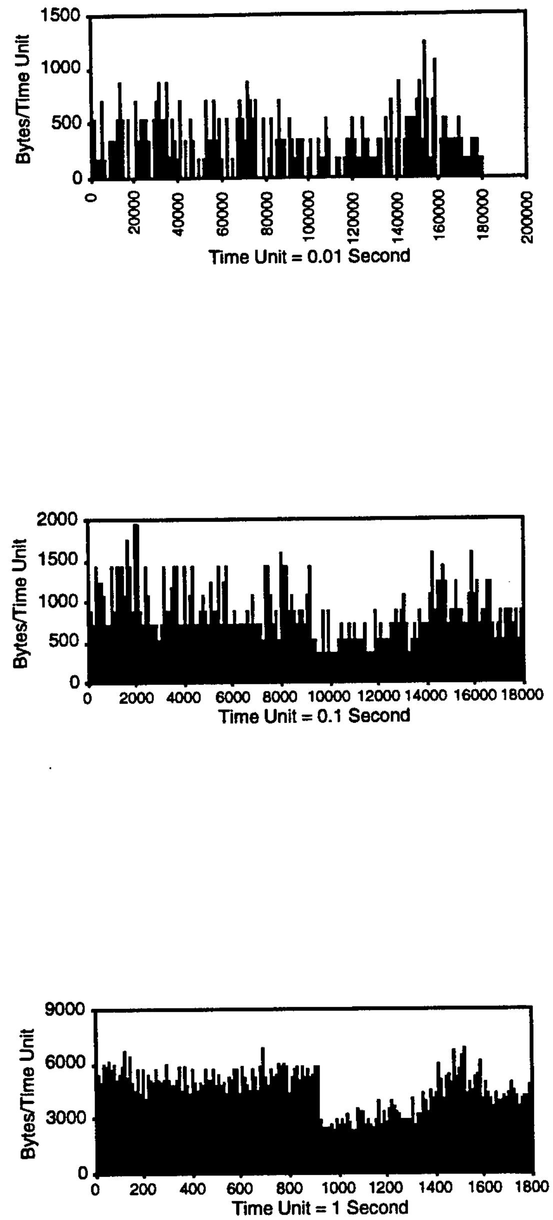

Fig. 15 Trace LAN2 at 3 Different Time Scales: 0.01 - 1 Second. The traffic shows non-diminishing burstiness in all three time scales.

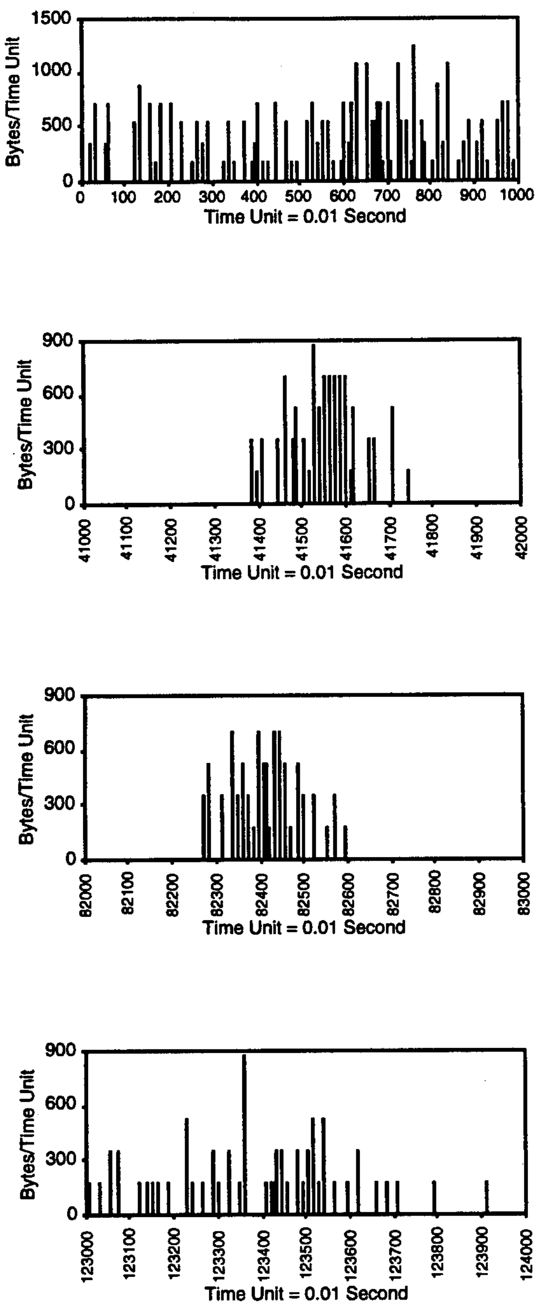

Fig. 16 Trace LAN2 in Four Non-Overlapping Time Windows. Each window contains 1000 0.01-second samples. The traffic forms clusters of approximately equal magnitudes (~600 bytes) in some windows.

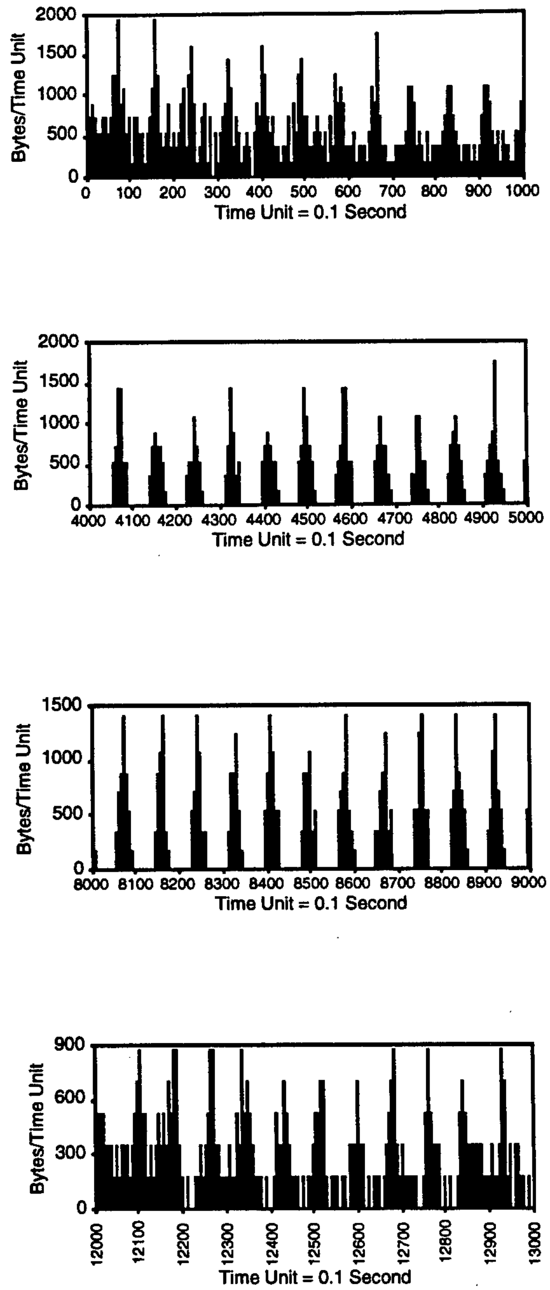

Fig. 17 Trace LAN2 in Four Non-Overlapping Time Windows. Each window contains 1000 0.1-second samples. The traffic forms clusters of approximately equal magnitudes (~1500 bytes) in some windows.

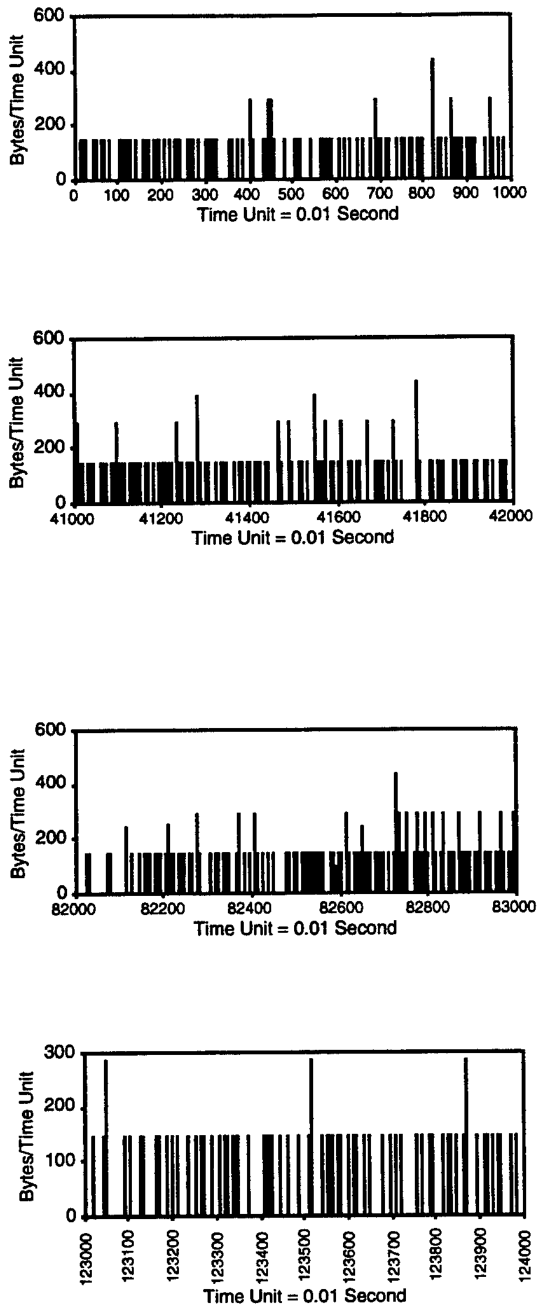

Fig. 18 Trace WAN2 in Four Non-Overlapping Time Windows. Each window contains 1000 0.01-second samples. The WAN traffic is more uniform than the LAN traffic as a result of the seven BRT algorithms housed in the AG. The WAN load level is uniform at about 150 bytes per 0.01 second.

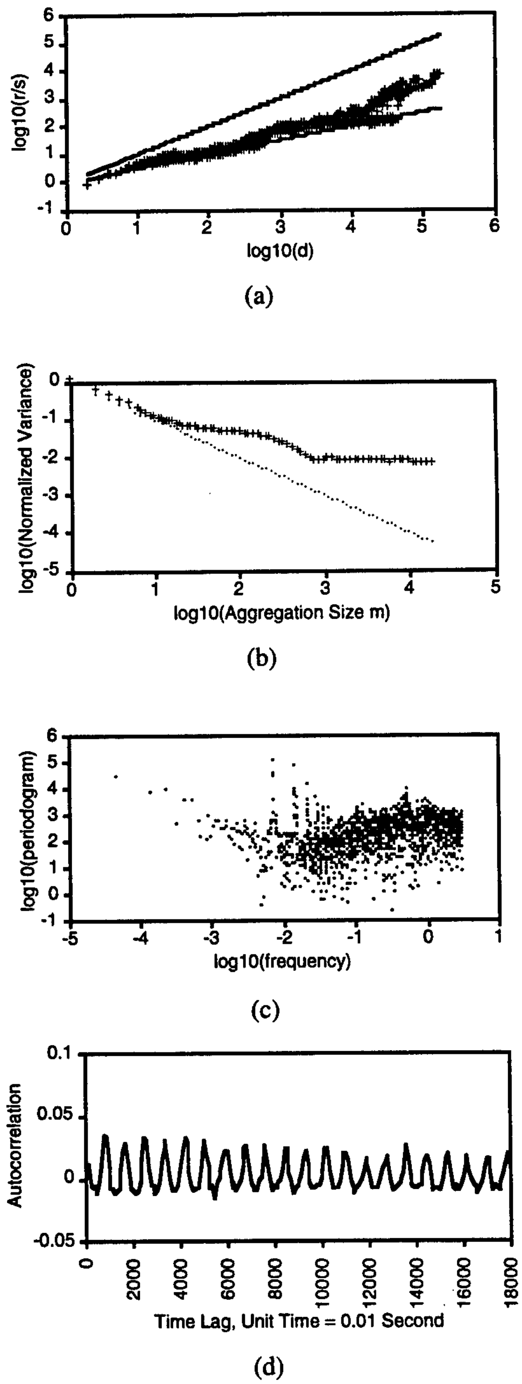

Fig. 19 Graphical Statistics for Trace LAN2. (a) *R/S* Plot: *d* denotes lag values or window sizes. The reference straight line of slope 1 (0.5 resp.) corresponds to H=1 (0.5 resp.). (b) Variance-Time Plot: The reference straight line of slope -1 corresponds to H=0.5. (c) Periodogram Plot. (d) Autocorrelation Plot.

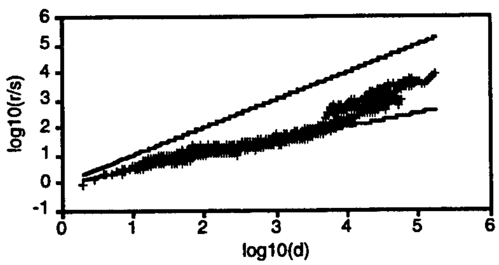

(a)

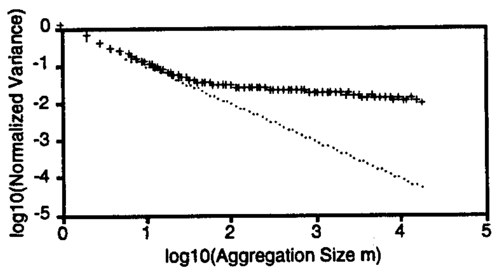

(b)

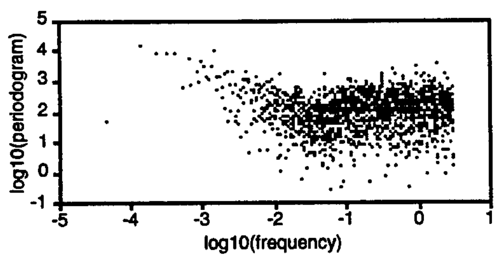

(c)

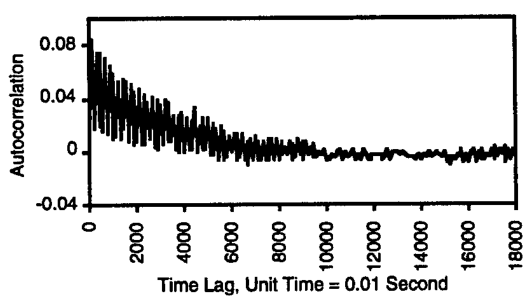

(d)

Fig. 20 Graphical Statistics for Trace WAN2. (a) R/S Plot: *d* denotes lag values or window sizes. The reference straight line of slope 1 (0.5 resp.) corresponds to H=1 (0.5 resp.). (b) Variance-Time Plot: The reference straight line of slope -1 corresponds to H=0.5. (c) Periodogram Plot. (d) Autocorrelation Plot.

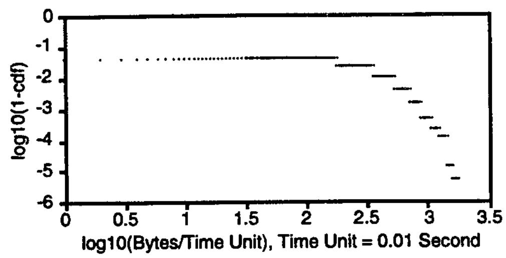

Fig. 21 Complementary Cumulative Distribution of Trace LAN2.The right tail varies somewhat as a straight line, which agrees with Pareto characteristics.

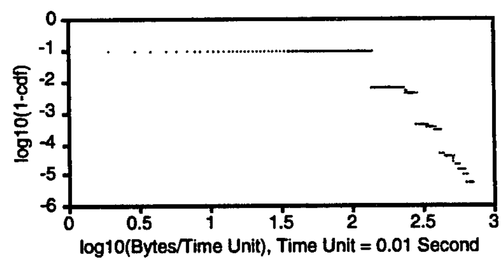

Fig. 22 Complementary Cumulative Distribution of Trace WAN2.The right tail varies somewhat as a straight line, which agrees with Pareto characteristics.

Appendix B

ADS	Advanced Distributed Simulation
AG	Application Gateway
ARPA	Advanced Research Projects Agency
AVTB	Aviation Test Bed
BRT	Bandwidth-Demand Reduction Technique
DARPA	Defense Advanced Research Projects Agency
Dlogger	Data Logger/LAN Logger
DIS	Distributed Interactive Simulation
DSI	Defense Simulation Internet
GMT	Greenwich Mean Time
HPAG	High Performance Application Gateway
IP	Internet Protocol
LAN	Local Area Network
LRD	Long-Range Dependence/Dependent
NRaD	Naval Command Control and Ocean Surveillance Center, Research, Development, Test and Evaluation Division
PDU	Protocol Data Unit
PICA	Protocol Independent Compression Algorithm
R/S	Rescaled Adjusted Range
SRD	Short-Range Dependence/Dependent
SIMNET	Simulation Network
STOW-E	Synthetic Theater of War-Europe
T1	Digital Signal 1
UDP	User Datagram Protocol
VBR	Variable Bit Rate
WAN	Wide Area Network
WANLogger	Wide Area Network Logger
WSS	Wide-Sense Stationary/Stationarity

Acknowledgment

We are grateful to Craig M. Barnhart for his helpful suggestions.

References

[Cox84] D. R. Cox, "Long-Range Dependence: A Review", in *Statistics: An Appraisal*, H. A. David and H. T. David (Eds), The Iowa State University Press, Ames, Iowa, 55-74, 1984

[CSTH95] J. Calvin, J. Seeger, G. Troxel, and D. Van Hook, "STOW Real Time Information Transfer and Networking System Architecture", preprint, 1995

[DLOR95] N. Duffield, J. Lewis, N. O'Connell, R. Russell, and F. Toomey, "Predicting Quality of Service for Traffic with Long-Range Fluctuations", *IEEE ICC*, 1995, 473-477

[GaWi94] M. Garrett and W. Willinger, "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic", *ACM SigComm*, Sept. 1994, 269-280

[HCNF94] D. Van Hook, J. Calvin, M. Newton, and D. Fusco, "An Approach to DIS Scaleability", *Eleventh DIS Workshop*, Orlando, Florida, September 26-30, 1994, 347-355

[HDLK95] C. Huang, M. Devetsikiotis, I. Lambadaris, and R. Kaye, "Fast Simulation for Self-Similar Traffic in ATM Networks", *IEEE ICC*, 1995, 438-444

[JaRo86] R. Jain and S. A. Routhier, "Packet Trains: Measurements and a New Model for Computer Network Traffic", *IEEE Journal on Selected Areas in Communications* **4**, 986-995, 1986

[Kana91] C. Kanarick, "A Technical Overview and History of the SIMNET Project", Bolt Beranek and Newman Inc., 1991, 104-111

[Klem74] V. Klemes, "The Hurst Phenomenon: a Puzzle?", *Water Resour. Res.*, **10**, 675-688

[LTWW94] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)", *IEEE/ACM Transactions on Networking*, **2**(1), 1-15, Feb. 1994

[NgBa95] G. Nguyen and S. Batsell, "Statistics for STOW-E Traffic", *NRL Technical Report*, Sept. 1995

[NRaD95] NRaD, "STOW-E Technical Analysis", Draft, 1995

[PaFl94] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling", *Technical report LBL*-35238, Lawrence Berkeley Laboratory, August 1994

[Paxs95] V. Paxson "Fast Approximation of Self-Similar Network Traffic", *Technical report LBL*-36750, Lawrence Berkeley Laboratory, April 1995

[SlLe95] B. Slimane and T. Le-Ngoc, "A Doubly Stochastic Poisson Model for Self-Similar Traffic", *IEEE ICC*, 1995, 456-460